Sublinear scaling of country attractiveness observed from Flickr dataset

Iva Bojic
SENSEable City Lab, MIT
Cambridge, MA, USA
Email: ivabojic@mit.edu

Ivana Nizetic-Kosovic
FER, University of Zagreb
Zagreb, Croatia
Email: ivana.nizetic@fer.hr

Alexander Belyi
SENSEable City Lab, SMART Centre
Singapore
Email: alex.bely@smart.mit.edu

Stanislav Sobolevsky
CUSP, New York University
Brooklyn, NY, USA
Email: sobolevsky@nyu.edu

Vedran Podobnik
FER, University of Zagreb
Zagreb, Croatia
Email: vedran.podobnik@fer.hr

Carlo Ratti
SENSEable City Lab, MIT
Cambridge, MA, USA
Email: ratti@mit.edu

arXiv:1601.02306v1 [cs.SI] 11 Jan 2016

***Abstract*—The number of people who decide to share their photographs publicly increases every day, consequently making available new almost real-time insights of human behavior while traveling. Rather than having this statistic once a month or yearly, urban planners and touristic workers now can make decisions almost simultaneously with the emergence of new events. Moreover, these datasets can be used not only to compare how popular different touristic places are, but also predict how popular they should be taking into an account their characteristics. In this paper we investigate how country attractiveness scales with its population and size using number of foreign users taking photographs, which is observed from Flickr dataset, as a proxy for attractiveness. The results showed two things: to a certain extent country attractiveness scales with population, but does not with its size; and unlike in case of Spanish cities, country attractiveness scales sublinearly with population, and not superlinearly.**

***Keywords*-human mobility; urban attraction; big data.**

I. Introduction

With the emergence of online social platforms (e.g., Twitter, Flickr, Facebook) not only do rich and famous people get attention, but also ordinary people can choose to participate in making their lives public. Moreover, even when people try to keep to themselves, their less concerned family members or friends can reveal a lot of things about their personal lives. Although we are still not fully aware to what extent will *living online* affect/affects us, the huge data produced by humans in digital world represents almost an inexhaustible resource for different scientific studies (e.g., land use classification methods [1] or determining economical potential of cities [2]).

The focus of this study is to investigate what catches people's eye using Flickr dataset of publicly available photographs and videos [3]. In our previous work we showed that city attractiveness in Spain scales superlinearly (exponent 1.5) with the city size [4] where this high exponent denotes that for example attractiveness of a city three times bigger than the other one, is expected not to be three times, but on average five times higher. Moreover, this finding seems to be robust across different city definitions and also is true for two other datasets: Twitter and dataset of bank card transaction records. In this paper we extend our study to investigating how attractiveness scales on a country level.

Various scaling studies have already been conducted on a city level as cities are important places that can transform a human life. While studies on socioeconomic parameters mostly show superlinear scaling laws, meaning that those parameters grow much faster than the city size, other studies on urban infrastructure reveal a sublinear relation to the city size. In that sense it has been shown that a bigger city boosts up human activity such as intensity of interactions [5], creativity [6], economic efficiency [7], as well as certain negative aspects: crime rates [6] or infectious diseases [8]. On the contrary to the socioeconomic parameters, urban infrastructure shows sublinear scaling proving that cities are mostly built in an efficient way [7], [8].

The novelty of our approach is not only in methods that we propose, but also in dataset that we use. In this paper we use publicly available Yahoo Flickr Creative Commons dataset, which is explained in more details in the next section, as a proxy for country attractiveness. Although this dataset was made publicly available only in 2014, scholars have already use it for building a user recommendation system for personalized tour based on their interests and points of interest visit duration [9] or summarizing real-time multimedia events using Flickr as a source for social media [10]. However, to the best of our knowledge, nobody has use it to investigate country attractiveness seen through the lens of people taking photographs and videos.

The rest of the paper is organized as follows. Section II introduces Yahoo Flickr Creative Commons dataset of more than 100 million media objects (i.e., photographs and videos). Section III first shows results of aggregated country attractiveness scaling with country population and its size for countries in the whole world and then zooms in to countries grouped in different regions. Finally, Section IV concludes the paper and gives guidelines for future work.

II. Dataset

In this study we use Yahoo Flickr Creative Commons dataset that is publicly available and can be downloaded from the following link: [11]. This dataset is the largest public multimedia collection that has ever been released and was created as a part of the Yahoo Webscope program. It contains 100 million media objects, which have been uploaded to Flickr between 2004 and 2014 and published under a CC commercial or non-commercial license, of which approximately 99.2 million are photographs. Each object in the dataset is represented by its metadata: object identifier, user identifier, time stamp when it was taken, location (i.e., latitude and longitudinal coordinates) where it was taken (if available), and CC license it was published under. Additionally, the metadata for some objects also contains object title, user tags, machine tags and description, as well as a direct link for downloading the content.

Since the focus of the paper is to investigate how country attractiveness scales with the country size, we had to perform reverse geocoding so that every location (i.e., latitude and longitude coordinates) where media object was created is translated into a human readable format (i.e., country). But before applying a reverse geocoding method to the given dataset, we had to prune the data. In data pruning process we first omitted objects that were not geo-tagged (i.e., more than 50%) and then objects that came with the wrong date format (652 in total). After both pruning and reverse geocoding we were left with 43,821,559 objects created in 242 countries around the world. Figure 1 shows unequal distribution of unique users and media objects per country.

However, the absolute numbers of media objects and unique visitors cannot be compared across different countries as countries are different in size and number of residents. When normalizing number of unique users with the number of people living in the country[1], we get completely different order of first five countries starting with Iceland where one Flicker user comes on 180 of its residents, followed by Greenland, San Marino, Andorra and Cayman Islands. With the exception of Iceland, this normalization shows a bias towards small countries that are touristically attractive. Moreover, when normalizing the number of media objects with the country size[2] we get similar order of the first five countries: Iceland, Greenland, Cayman Islands, United States Virgin Islands and Palau.

[1]Information about world population (i.e., number of people living in each country) was downloaded from the world bank website (http://data.worldbank.org/indicator/SP.POP.TOTL). Since Flickr dataset includes media objects created between years 2004 and 2014, in our normalization process we do not use the population for a specific year, but calculate the average number of people living in the country for the same period of ten years.

[2]Country's total area was also downloaded from the world bank website (http://data.worldbank.org/indicator/AG.LND.TOTL.K2) and was prepared in the same way as population, i.e., calculating an average value for the time period between 2004 and 2014.

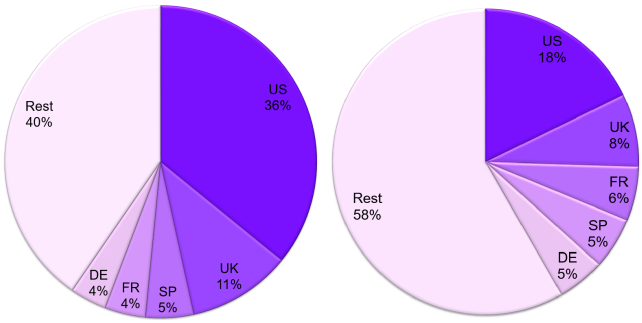

Figure 1. a) Distribution of media objects per country, b) distribution of unique users per country. Objects and users are not equally distributed since 60% of media objects and more than 40% of unique visitors belong to only five countries (i.e., US, UK, SP, DE, FR).

Finally, Figure 2 shows an average number of media objects that every user produces in a given country ranging from more than 200 objects per user made in Taiwan and the United States to less than 10 made in for example Chad and San Marino.

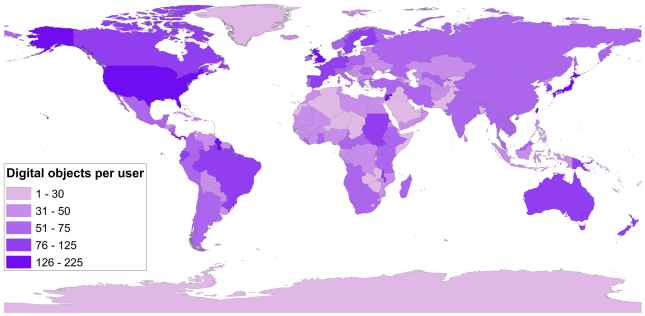

Figure 2. Map showing an average number of media objects produced by users in a specific country. The first five countries are as follows: Taiwan, United States, Panama, Guyana, Japan with 222, 215, 197, 174 and 162 media objects per user respectively.

III. Scaling of country attractiveness

Rather than including all human activity captured in the dataset, for the purpose of investigating country attractiveness scaling laws, we only included foreign users and their activity. Since Flickr dataset does not include information about user home countries, we used the same methodology as we used in our previous papers [12], [13] to infer where people live. We say that a particular user lives in a particular country if there he/she made the maximum number of media objects and spend the maximal number of days. Then the activity of those users for whom we detected their home countries is used when calculating aggregated country attractiveness in two different ways: number of media objects taken by foreign users compared to the country size and its population. Aggregated country attractiveness denotes that we included all media objects without taking into consideration time when they were made.

Figure 3 shows the results when fitting a power-law dependence $A \sim ap^{\beta}$ or $A \sim as^{\beta}$ to the calculated number of media objects denoting aggregated country attractiveness A and population p or size s. Fitting is performed on a log-log scale where it becomes a simple linear regression problem $log(A) \sim log(a) + \beta \cdot log(p)$ or $log(A) \sim log(a) + \beta \cdot log(s)$. As proposed in [8] scaling can fall in three different categories depending on the value of parameter β: linear ($\beta = 1$), sublinear ($\beta < 1$) and superlinear ($\beta > 1$). Unlike city attractiveness [4], country attractiveness for countries all around the world scales sublinearly with both country population p and size s. A measure of goodness-of-fit of linear regression R^2 is 26.7% for population p and 11.9% for size s, denoting that sublinear scaling trend of country attractiveness with the population is more emphasized than scaling trend with country size. However, values for both R^2 are small enough meaning that we cannot get a perfect fit for all counties in the world. Therefore, we divided world countries into eleven regions[3] and checked if scaling laws are more present in certain regions rather than all world.

Table I shows that all countries belonging to different regions are sublinearly scaling with their population. Regions are sorted descending by values of R^2, i.e., best fit. Results show that almost a perfect fit (almost 99%) is for region called Northern America with only four countries, followed by Western Europe which consists of 25 countries. However, the number of countries in the region is not proportional with a measure of goodness-of-fit of linear regression R^2 as regions Asia (ex. Near East) and Near East have the same number of countries as Western Europe or less and their fit is less than 26%. Aggregated country attractiveness for Western Europe and Asia (ex. Near East), which both are composed of 25 countries, is shown in Figure 4.

[3]Data used in this study can be downloaded from the following link: http://www.statvision.com/webinars/Countries%20of%20the%20world.xls.

Table I
WORLD REGIONS WITH NUMBER OF COUNTRIES BELONGING TO THEM, POPULATION, SCALING COEFFICIENT β AND PARAMETER R^2.

Region name (no of countries)	Population	β	R^2 (%)
Northern America (4)	340,100,983	0.777	98.5
Western Europe (25)	409,402,265	0.715	93.8
Baltics (3)	6,607,817	-0.499	89.0
Northern Africa (6)	161,701,943	0.964	75.4
Commonwealth of Independent States (12)	280,892,606	0.933	63.0
Oceania (16)	35,557,584	0.850	59.2
Latin America and Caribbean (40)	589,003,980	0.466	57.6
Eastern Europe (12)	114,521,789	0.778	40.7
Sub-Saharan Africa (48)	836,231,137	0.638	27.6
Asia (ex. Near East) (25)	3,787,575,787	0.363	25.7
Near East (14)	205,611,734	0.344	11.0

Since there were not obvious correlations between R^2 and parameters form Table I, we calculated the correlation coefficient for R^2 and country area, population, GDP, density, coastline and urban population. The correlation coefficient, which is a value between -1 and $+1$, tells how strongly two variables are related to each other. The smallest correlation coefficient was between R^2 and density (-0.102), while the biggest was with GDP (0.678).

IV. Conclusions

In this study we analyzed aggregated country attractiveness in dependence with its population and size for countries around the world. The results on country attractiveness showed that there is not a direct correlation between it and country size, unlike with population where different regions of the world showed different behavior. Namely, regions like Western Europe, Northern America and Baltics showed linearity on log-log scale, while regions like Sub-Saharan Africa, Asia and Near East are not well fitted.

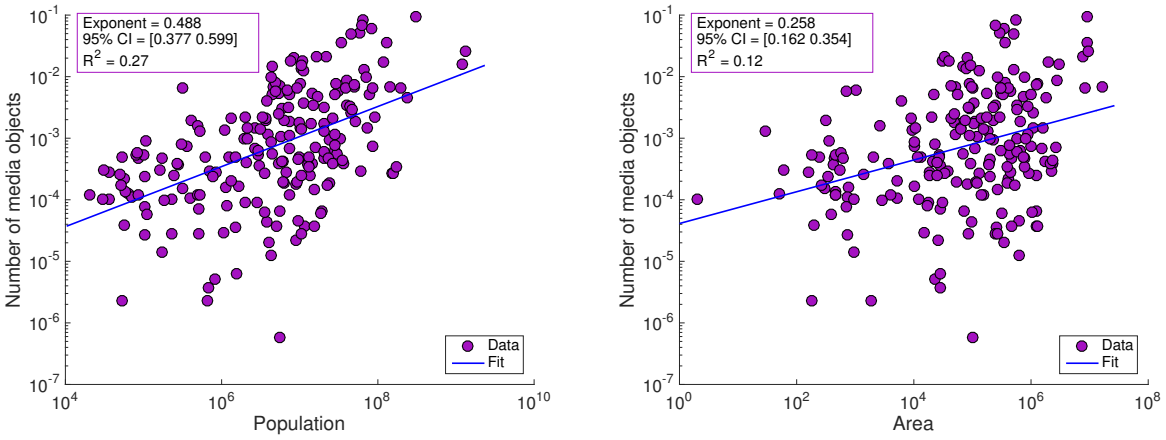

Figure 3. Scaling of aggregated county attractiveness where X-axis represents the number of people living in the country or the country size, while Y-axis represents the fraction of the number of media objects made in a particular country versus the total amount of objects made in all countries.

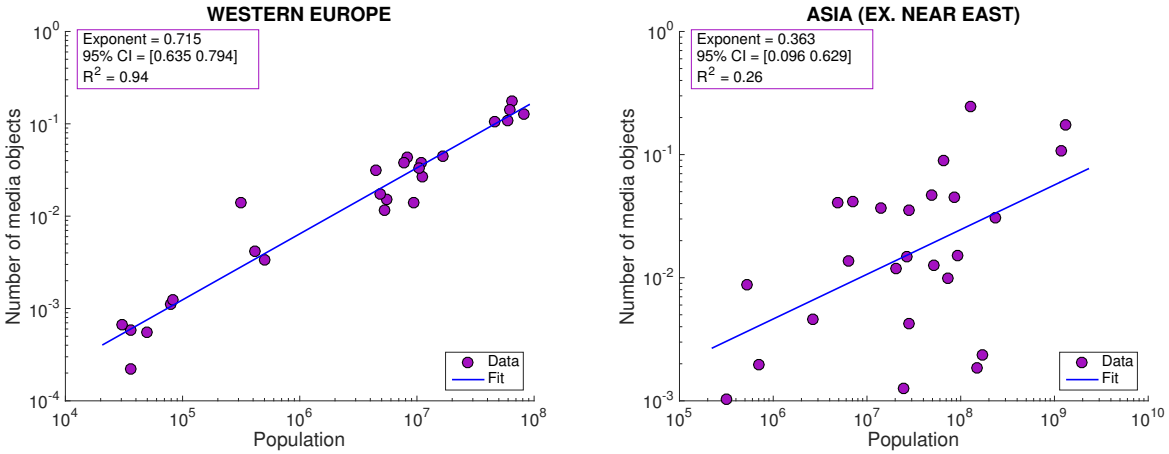

Figure 4. Scaling of aggregated county attractiveness for Western Europe (on left) and Asia excluding Near East (on right).

Further, we also investigated correlations between the goodness-of-fit R^2 and various parameters of the countries (e.g., density, coastline) and discovered the most significant correlation with country GDP, i.e., 0.678. Regardless to the goodness-of-fit, all regions showed sublinear scaling with their population, rather than superlinear one discovered within prior studies on aggregated city attractiveness.

ACKNOWLEDGMENTS

The authors would like to thank BBVA, MIT SMART Program, Center for Complex Engineering Systems (CCES) at KACST and MIT, Accenture, Air Liquide, The Coca Cola Company, Emirates Integrated Telecommunications Company, The ENELfoundation, Ericsson, Expo 2015, Ferrovial, Liberty Mutual, The Regional Municipality of Wood Buffalo, Volkswagen Electronics Research Lab, UBER, and all the members of the MIT Senseable City Lab Consortium for supporting the research. Part of this research was also funded by the research project "Managing Trust and Coordinating Interactions in Smart Networks of People, Machines and Organizations", funded by the Croatian Science Foundation.

REFERENCES

[1] S. Grauwin, S. Sobolevsky, S. Moritz, I. Gódor, and C. Ratti, "Towards a comparative science of cities: Using mobile traffic records in New York, London and Hong Kong," in *Computational Approaches for Urban Environments*, ser. Geotechnologies and the Environment, 2015, vol. 13, pp. 363–387.

[2] S. Sobolevsky, E. Massaro, I. Bojic, J. M. Arias, and C. Ratti, "Predicting regional economic indices using big data of individual bank card transactions," *arXiv preprint arXiv:1506.00036*, 2015.

[3] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li, "The new data and new challenges in multimedia research," *arXiv preprint arXiv:1503.01817*, 2015.

[4] S. Sobolevsky, I. Bojic, A. Belyi, I. Sitko, B. Hawelka, J. M. Arias, and C. Ratti, "Scaling of city attractiveness for foreign visitors through big data of human economical and social media activity," *arXiv preprint arXiv:1504.06003*, 2015.

[5] M. Schläpfer, L. Bettencourt, S. Grauwin, M. Raschke, R. Claxton, Z. Smoreda, G. B. West, and C. Ratti, "The scaling of human interactions with city size," *Journal of the Royal Society Interface*, vol. 11, no. 98, pp. 1–9, 2014.

[6] L. M. Bettencourt, J. Lobo, D. Strumsky, and G. B. West, "Urban scaling and its deviations: Revealing the structure of wealth, innovation and crime across cities," *PLoS One*, vol. 5, no. 11, pp. 1–9, 2010.

[7] L. M. Bettencourt, "The origins of scaling in cities," *Science*, vol. 340, no. 6139, pp. 1438–1441, 2013.

[8] L. M. Bettencourt, J. Lobo, D. Helbing, C. Kühnert, and G. B. West, "Growth, innovation, scaling, and the pace of life in cities," *Proceedings of the National Academy of Sciences*, vol. 104, no. 17, pp. 7301–7306, 2007.

[9] K. H. Lim, J. Chan, C. Leckie, and S. Karunasekera, "Personalized tour recommendation based on user interests and points of interest visit durations," *Under Submission*, 2015.

[10] R. R. Shah, A. D. Shaikh, Y. Yu, W. Geng, R. Zimmermann, and G. Wu, "Eventbuilder: Real-time multimedia event summarization by visualizing social media," in *Proceedings of the 23rd ACM International Conference on Multimedia*, 2015, pp. 185–188.

[11] Yahoo! Webscope dataset YFCC-100M. [Online]. Available: http://labs.yahoo.com/Academic-Relations

[12] S. Paldino, I. Bojic, S. Sobolevsky, C. Ratti, and M. C. González, "Urban magnetism through the lens of geo-tagged photography," *EPJ Data Science*, vol. 4, no. 1, pp. 1–17, 2015.

[13] I. Bojic, E. Massaro, A. Belyi, S. Sobolevsky, and C. Ratti, "Choosing the right home location definition method for the given dataset," *arXiv preprint arXiv:1510.03715*, 2015.